

Filters: when, why, and how (not) to use them

Alain de Cheveigné (1, 2, 3), Israel Nelken (4)

(1) Laboratoire des Systèmes Perceptifs, UMR 8248, CNRS, France.

(2) Département d'Etudes Cognitives, Ecole Normale Supérieure, PSL, France.

(3) UCL Ear Institute, United Kingdom.

(4) Edmond and Lily Safra Center for Brain Sciences and the Silberman Institute of Life Sciences, Hebrew University Jerusalem, Israel.

CORRESPONDING AUTHOR:

Alain de Cheveigné, Audition, DEC, ENS, 29 rue d'Ulm, 75230, Paris, France,
Alain.de.Cheveigne@ens.fr, phone 0033144322672, 00447912504027.

approximate number of characters (with spaces): 79000

1 Summary

2 Filters are commonly used to reduce noise and improve data quality. Filter theory
3 is part of a scientist’s training, yet the impact of filters on interpreting data is
4 not always fully appreciated. This paper reviews the issue, explaining what is a
5 filter, what problems are to be expected when using them, how to choose the right
6 filter, or how to avoid filtering by using alternative tools. Time-frequency analysis
7 shares some of the same problems that filters have, particularly in the case of
8 wavelet transforms. We recommend reporting filter characteristics with sufficient
9 details, including a plot of the impulse or step response as an inset.

10 1 Introduction

11 One of the major challenges of brain science is that measurements are contami-
12 nated by noise and artifacts. These may include environmental noise, instrumental
13 noise, or signal sources within the body that are not of interest in the context of
14 the experiment (“physiological noise”). The presence of noise can mask the target
15 signal, or interfere with its analysis. However, if signal and interference occupy
16 different spectral regions, it may be possible to improve the signal-to-noise ratio
17 (SNR) by applying a *filter* to the data.

18 For example, a DC component or slow fluctuation may be removed with a
19 high-pass filter, power line components may be attenuated by a notch filter at
20 50 Hz or 60 Hz, and unwanted high-frequency components may be removed by
21 “smoothing” the data with a low-pass filter. Filtering takes advantage of the dif-
22 ference between spectra of noise and target to improve SNR, attenuating the data
23 more in the spectral regions dominated by noise, and less in those dominated by
24 the target.

25 Filters are found at many stages along the measurement-to-publication pipeline
26 (Fig. 1). The measuring rig or amplifier may include a high-pass filter and pos-
27 sibly a notch filter, the analog-to-digital (AD) converter is preceded by a low-
28 pass antialiasing filter, preprocessing may rely on some combination of high-pass,
29 lowpass and notch filters, data analysis may include bandpass or time-frequency
30 analysis, and so on. Filters are ubiquitous in brain data measurement and analysis.

31 The improvement in SNR offered by the filter is welcome, but filtering affects

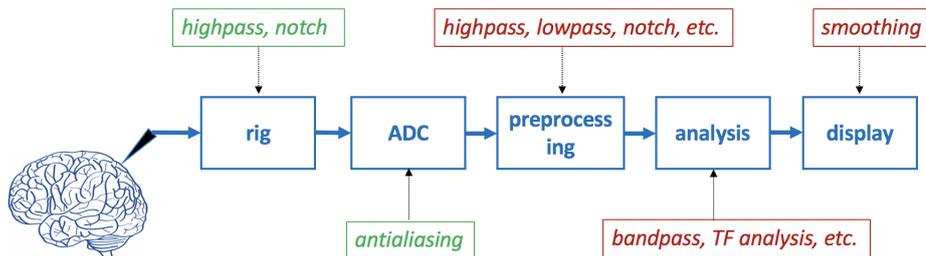


Figure 1: A typical recording-to-publication data pipeline, showing where filters are applied. Filters are analog in the first stages (green) and digital in subsequent stages (red). The recording rig might include a high-pass filter (implicit in the case of AC coupling), and perhaps also a notch filter to attenuate line frequency power. The analog-to-digital converter is preceded by a low-pass antialiasing filter. In data preprocessing it is common to apply a high-pass filter to remove slow drift components, and a low-pass filter to attenuate noise (often spread over the entire spectrum), or to avoid antialiasing when the data are down-sampled. Data analysis might involve bandpass filtering (for example to isolate a standard frequency band such as "alpha" or "gamma") or time-frequency analysis. Data display or plotting might call for additional smoothing (low-pass filtering).

32 also the target signal in ways that are sometimes surprising. Obviously, any com-
 33 ponents of the target signal that fall within the stop band of the filter are lost. For
 34 example applying a 50 Hz notch filter to remove power line artifact might also re-
 35 move brain activity within the 50 Hz region. The experimenter who blindly relies
 36 on the filtered signal is blind to features suppressed by the filter.

37 Harder to appreciate are the distortions undergone by the target. Such dis-
 38 tortions depend on the frequency characteristics of the filter, including both am-
 39 plitude and phase characteristics (which are often not reported). The output of a
 40 filter is obtained by *convolution* of its input with the impulse response of the filter,
 41 which is a fancy way of saying that each sample of the output is a weighted sum
 42 of several samples of the input. Each sample therefore depends on a whole seg-
 43 ment of the input, spread over time. Temporal features of the input are smeared in
 44 the output, and conversely new "features" may appear in the output that were not
 45 present in the input to the filter.

46 We first explain what is a filter in detail, and how filters are involved in data
 47 analysis. Then we review the main issues that can arise, and make suggestions on
 48 how to fix them. Importantly, similar issues occur also in time-frequency analy-

49 ses, such as spectrograms and wavelet transforms, which are based on a collection
50 of filters (a filterbank). Finally, we list a number of recommendations that may
51 help investigators identify and minimize issues related to the use of filters, and we
52 suggest ways to report them so that readers can make the best use of the informa-
53 tion that they read. In this paper, “filter” refers to the familiar one-dimensional
54 convolutional filter (e.g. high-pass or band-pass) applicable to a single-channel
55 waveform, as opposed to “spatial filters” applicable to multichannel data.

56 **2 What is a filter?**

57 For many of us, a filter is “a thing that modifies the spectral content of a signal”.
58 For the purposes of this paper, however, we need something more precise. A filter
59 is an operation that produces each sample of the output waveform y as a weighted
60 sum of several samples of the input waveform x . For a digital filter:

$$y(t) = \sum_{n=0}^N h(n)x(t - n) \quad (1)$$

61 where t is the analysis point in time, and $h(n), n = 0, \dots, N$ is the impulse re-
62 sponse. This operation is called convolution.

63 We expect the reader to fall into one of three categories: (a) those who under-
64 stand and feel comfortable with this definition, (b) those who mentally transpose it
65 to the frequency domain where they feel more comfortable, and (c) those who re-
66 main mind-boggled. Categories (b) and (c) both need assistance, and that is what
67 this section is about. Category (b) need assistance because a frequency-domain
68 account is incomplete unless phase is taken into account, but doing so is mentally
69 hard and often not so illuminating. It is often easier to reason in the time domain.

70 For the mind-boggled, the one important idea to retain is that every sample of
71 the output depends on *multiple* samples of the input, as illustrated in Fig. 2 (top).
72 Conversely, each sample $x(t)$ of the input impacts several samples $y(t + n)$ of
73 the output (Fig. 2, bottom). As a result, the signal being filtered is smeared along
74 the temporal axis, and temporal relations between filtered and original waveforms
75 are blurred. For example, the latency between a sensory stimulus and a brain
76 response, a straightforward notion, becomes less well defined when that brain
77 response is filtered.

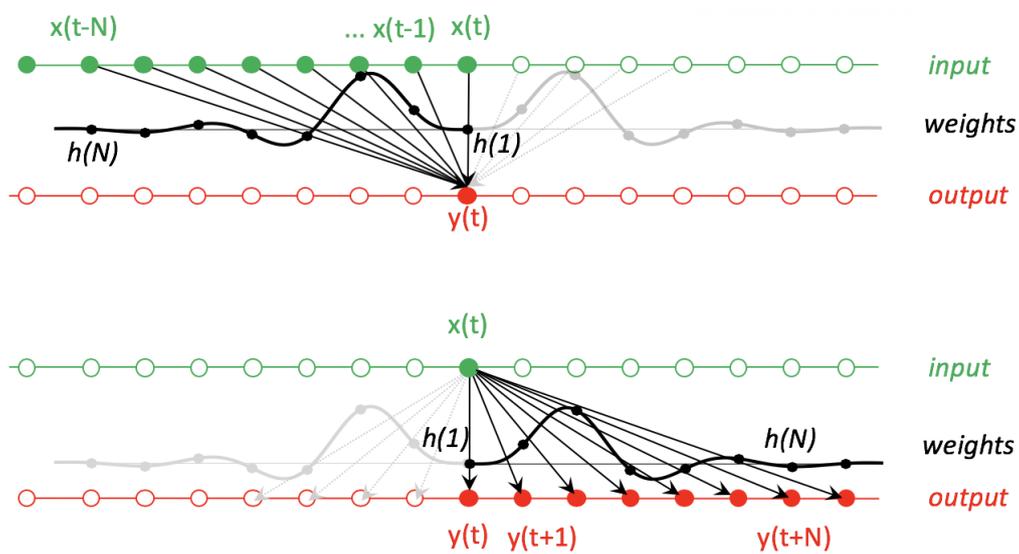


Figure 2: Filtering. Top: each sample of the output y is the sum of samples of the input x weighted by the impulse response h . For a causal filter, only past or present samples of the input make a contribution (black). For an acausal filter, future samples too can contribute (gray). Bottom: another way of describing this process is that each sample of the input x affects multiple samples of the output y , with a weight determined by the impulse response h .

78 The exact way in which the output of a filter differs from its input depends
79 upon the filter, i.e. the values $h(n)$ of the impulse response. Some filters may
80 smooth the input waveform, others may enhance fast variations. There is a con-
81 siderable body of theory, methods, and lore on how best to design and implement
82 a filter for the needs of an application.

83 Expert readers will add that a filter is a *linear system*, that $h(n)$ is not expected
84 to change over time (*linear time-invariant system*), that in addition to *causal* filters
85 described by Eq. 1, there are *acausal* filters for which the series $h(n)$ includes
86 also negative indices (gray lines in Fig. 2), that N may be finite (*finite impulse*
87 *response, FIR*) or infinite (*infinite impulse response, IIR*). IIR filters are often
88 modeled after familiar analogue filter designs (e.g. Butterworth).

89 Essentially everything we discuss below is true for these more general notions
90 of filtering. Expert readers will also recognize that Eq. 1 can be substituted by the
91 simpler equation $Y(\omega) = H(\omega)X(\omega)$ involving the Fourier transforms of $x(t)$,
92 $y(t)$ and $h(n)$, that neatly describes the effects of filtering in the frequency domain
93 as a product of two complex functions, the transfer function of the filter $H(\omega)$ and
94 the Fourier transform of the input, $X(\omega)$. The *magnitude transfer function* $|H(\omega)|$
95 quantifies the amount of attenuation at each frequency ω .

96 A special mention should be made of *acausal* filters. These are filters for
97 which each sample of the output depends also on future samples of the input, i.e.
98 we must modify Eq. 1 to include negative indices $n = -N', \dots, -1$. All physical
99 systems must be causal (the future cannot influence the past) so this filter cannot
100 represent a physical system, nor could it be implemented in a real time processing
101 device. However, for offline data analysis we can take samples from anywhere
102 in the data set, so in that context acausal filters are realizable. In particular, it is
103 common to use *zero phase* filters, for which the impulse response is symmetrical
104 relative to zero. The Matlab function `filtfilt` applies the same filter to the
105 data twice, forward and backward, effectively implementing a zero-phase filter.

106 While acausal filters are easy to apply, interpreting their output requires spe-
107 cial care. An important goal of neuroscience is to determine causal relations, for
108 example between a stimulus and brain activity, or between one brain event and
109 another, and we must take care that these relations are not confused by an acausal
110 stage in the data analysis.

111 For an *IIR* filter, the output depends on all samples from the start of the data,

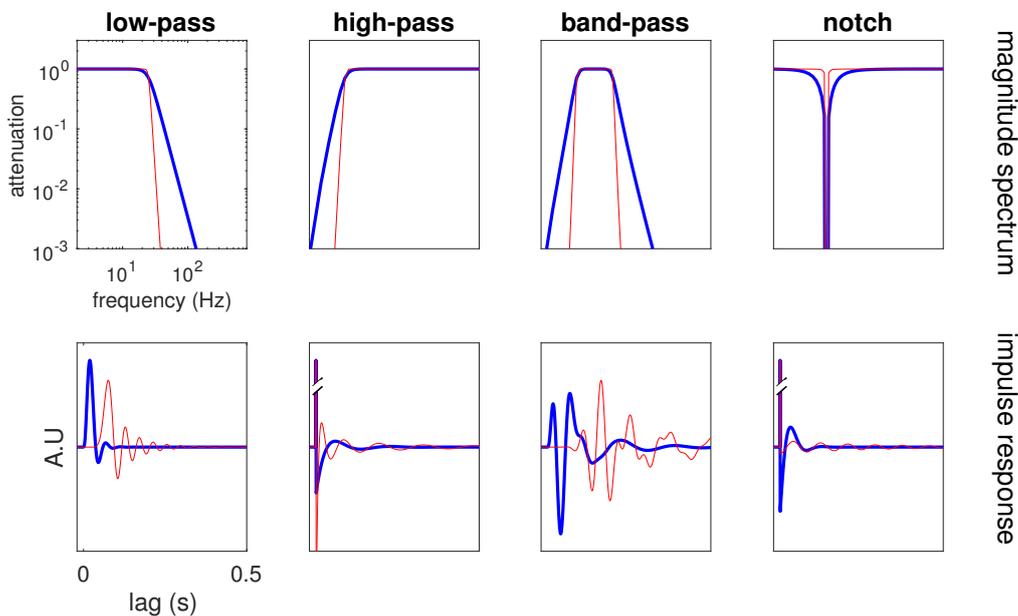


Figure 3: Typical magnitude transfer function shapes (top) and the associated impulse responses (bottom). The low-pass filter attenuates high frequencies, the high-pass attenuates low frequencies, the band-pass attenuates out-of band frequencies, the notch attenuates a narrow band of frequencies. The steeper the transition in the frequency domain, the more extended the impulse response (red). The steepness of the transitions depends on the type and order of the filter. Low-pass, high-pass and band-pass are Butterworth filters of order 4 and 16, notch filters are second-order filters with Q factors (ratio of bandwidth to centre frequency) of 1 and 10. Impulse responses for high-pass and notch include a high amplitude impulse, plotted here with a break.

112 previous samples being treated as 0. If the IIR filter is acausal it can also depend
 113 on all samples until the end of the data, samples beyond the end being treated
 114 as 0. For a filter implemented in the *Fourier domain*, each output sample $y(t)$
 115 depends potentially on all input samples $x(t)$ that are used to compute the Fourier
 116 transform, i.e. every sample within the analysis window.

117 Figure 3 illustrates four common types of filter: low-pass, high-pass, band-
 118 pass, and notch. The upper plots show the magnitude transfer function (on a
 119 log-log scale) and the bottom plots show the impulse response of each filter. For
 120 high-pass and notch filters, the impulse response includes a one-sample impulse
 121 ("Dirac") of amplitude much greater than the rest (plotted here using a split ordi-
 122 nate). For each filter two versions are shown, one with with shallow (blue) and the

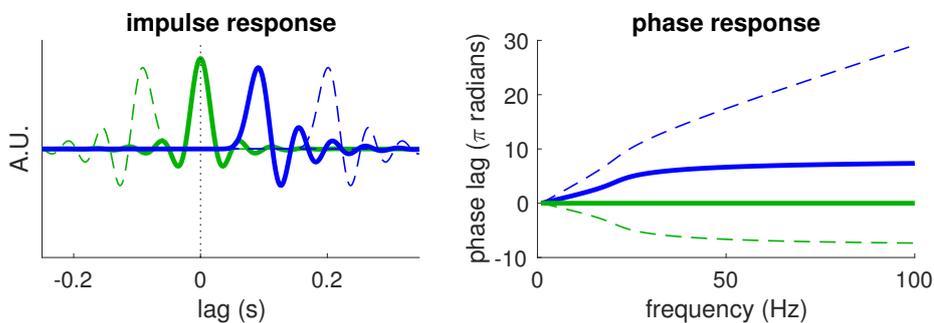


Figure 4: Left: impulse responses that all yield the same magnitude transfer function (lowpass, similar to that shown in Fig. 3 top left). The filters in blue are causal, those in green acausal. The filter in thick green is zero-phase. All examples are implemented as a cascade of two Butterworth lowpass filters of order 8 and cutoff 10 Hz. Thick blue: both impulse responses are convolved. Blue dashed: same but the result is delayed. Green dashed: same as thick blue but time-reversed. Thick green: one impulse response is time-reversed then convolved with the other. Right: corresponding phase responses. In subsequent figures, causal filters are plotted in blue, acausal in green.

123 other with steep (red) frequency transitions. Note that a filter with a steep transi-
 124 tion in the frequency domain tends to have an impulse response that is extended
 125 in the time domain.

126 Also important to note is that different impulse responses can yield the same
 127 magnitude transfer function. Figure 4 (left) shows four impulse responses that
 128 all share the same magnitude frequency characteristic (low-pass, similar to that
 129 shown in Fig. 3) but differ in their phase characteristics (plotted on the right).
 130 Magnitude and phase together fully specify a filter (as does the impulse response).
 131 Among all the filters that yield the same magnitude frequency response, one is
 132 remarkable in that it is causal and has *minimum phase* over all frequency (thick
 133 blue). Another is remarkable in that it has *zero phase* over all frequency (thick
 134 green). It is acausal.

135 3 Uses of filters

136 **Antialiasing.** Ubiquitous, if rarely noticed, is the hardware “antialiasing” filter
 137 that precedes analog-to-digital conversion within the measuring apparatus. Data
 138 processing nowadays is almost invariably done in the digital domain, and this re-

139 quires signals to be sampled at discrete points in time so as to be converted to a
140 digital representation. Only values at the sampling points are retained by the sam-
141 pling process, and thus the digital representation is ambiguous: the same set of
142 numbers might conceivably reflect a different raw signal. The ambiguity vanishes
143 if the raw signal obeys certain conditions, the best known of which is given by
144 the *sampling theorem*: if the original signal's spectrum contains no power beyond
145 the Nyquist frequency (one half the sampling rate) then it can be perfectly recon-
146 structed from the samples. The antialiasing filter aims to enforce this condition
147 ("Nyquist condition"). A hardware antialiasing filter is usually applied before
148 sampling, and a software antialiasing filter may later be applied if the sampled
149 data are downsampled or resampled.

150 **Smoothing / low-pass filtering.** Phenomena of interest often obey slow dynam-
151 ics. In that case, high-frequency variance can safely be attributed to irrelevant
152 noise fluctuations and attenuated by low-pass filtering. Smoothing is also often
153 used to make data plots visually more palatable, or to give more emphasis on
154 longer-term trends than on fine details.

155 **High-pass filtering to remove drift and trends.** Some recording modalities
156 such as electroencephalography (EEG) or magnetoencephalography (MEG) are
157 susceptible to DC shifts and slow drift potentials or fields, upon which ride the
158 faster signals of interest (Huigen et al., 2002; Kappenman and Luck, 2010; Van-
159 hatalo et al., 2005). Likewise, in extracellular recordings, spikes of single neurons
160 ride on slower events, such as negative deflections of the local field potential (LFP)
161 that often precede spikes, or the larger and slower drifts due to the development of
162 junction potentials between the electrode tip and the brain tissue. High-pass filters
163 are the standard tool to remove such slow components prior to data analysis. A
164 hardware high-pass filter might also be included in the measurement apparatus to
165 remove DC components prior to conversion so as to make best use of the lim-
166 ited range of the digital representation. This is the meaning of "ac coupling" on
167 an oscilloscope - it consists of the application of a high-pass filter to the signal.
168 Amplifiers for recording extracellular brain activity are usually AC coupled.

169 **Notch filtering.** Electrophysiology is often plagued with power line noise (50
170 or 60 Hz and harmonics) coupled electrically or magnetically with the recording
171 circuits. While it is best eliminated at its source by careful equipment design
172 and shielding, this is not always successful, nor is it applicable to data already
173 gathered. Notch filtering is often used to mitigate such power line noise.

174 **Band-pass filtering.** It has become traditional to interpret brain activity as com-
175 ing from frequency bands with names such as *alpha*, *beta*, *theta*, etc., and data
176 analysis often involves applying one or more band-pass filters to isolate particular
177 bands, although the consensus is incomplete as to the boundary frequencies or the
178 type of filter to apply.

179 **Time-frequency analysis.** One prominent application of filtering is *time-frequency*
180 (TF) analysis. A TF representation can be viewed as the time-varying magnitude
181 of the data at the outputs of a filterbank. A filterbank is an array of filters that dif-
182 fer over a range of parameter values (e.g. center frequencies and/or bandwidths).
183 The indices of the filters constitute the frequency axis, while the time series of
184 their output magnitude unfolds along the time axis of the TF representation. The
185 time-varying magnitude is obtained by applying a non-linear transform to the filter
186 output, such as half-wave rectification or squaring, possibly followed by a power
187 or logarithmic transform. The time-varying phase in each channel may also be
188 represented.

189 **4 How do filters affect brain data?**

190 The answer to this question depends on the data and on the filter. In this section
191 we review a number of archetypical “events” that might occur within a time series
192 of brain activity, and look at how they are affected by commonly used filters.

193 **Impulse or spike.** Brain events that are temporally localized, for example a neu-
194 ronal “spike”, can be modeled as one or a few impulses. It is obvious from Eq. 1
195 that such events must be less localized once filtered, as summarized schematically
196 in Fig. 5. The response is *spread over time*, implying that the temporal location of

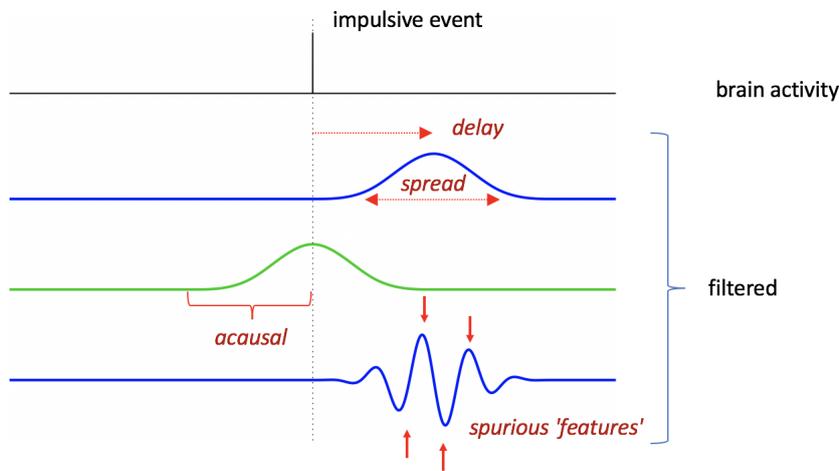


Figure 5: Effects of filtering on a temporally localized event (impulse or spike). The response is spread over time (i.e. no longer precisely localized in time), and delayed if the filter is causal. The overall delay is eliminated if the filter is zero-phase (green), but the response is then *acausal*, i.e. part of it occurs before the event. The response may include multiple spurious 'features' due to filter ringing.

197 the event is less well defined. It is *delayed* if the filter is causal. The delay may be
 198 avoided by choosing a zero-phase filter (green), but the response is then *acausal*.
 199 If the impulse response has multiple modes, these may appear misleadingly as
 200 multiple *spurious events*, confusing the analysis.

201 The nature and extent of these effects depends on the filter, and can be judged
 202 by looking at its impulse response. Figure 6 shows impulse responses of a selec-
 203 tion of commonly used filters (others were shown in Fig. 3). The left-hand plot
 204 shows the time course of the impulse response, and the right-hand plot displays
 205 the logarithm of its absolute value using a color scale, to better reveal the low-
 206 amplitude tail. The first three examples (A-C) correspond to low-pass filters with
 207 the same nominal cutoff (10 Hz). The next two (D-E) are low pass filters with
 208 nominal cutoff 20 Hz. The following three (F-H) are band-pass filters. Two of the
 209 filters are zero-phase (B and E, in green), the others are causal (blue).

210 The response of the first two filters is relatively short and unimodal, that of
 211 the others is more extended and includes excursions of both signs. The temporal
 212 span is greater for filters of high order (compare F and G) and for lower frequency
 213 parameters (compare C and D). Bandpass filters have relatively extended impulse

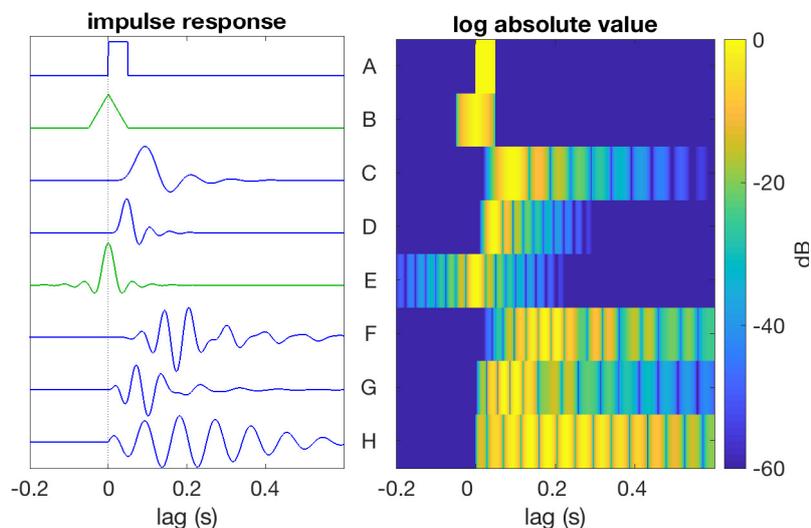


Figure 6: Impulse responses of typical filters. The left panel shows the impulse response time series, the right plot shows the log of its absolute value (truncated at -60 dB) on a color scale to reveal the low-amplitude tail. A: boxcar of duration 50 ms, B: same, applied using `filtfilt`, C: Butterworth lowpass, cutoff 10 Hz, order 8, D: same as C, cutoff 20 Hz, E: same as D, applied using `filtfilt`, F: Butterworth bandpass, 10-20 Hz, order 8, G: same, order 2, H: same as G, 10-12 Hz. Filters plotted in green are acausal.

214 responses, particularly if the band is narrow or the slopes of the transfer function
 215 steep.

216 Of course, real brain events differ from an infinitely narrow unipolar impulse,
 217 for example they have finite width, and the response to such events will thus differ
 218 somewhat from the ideal impulse response. As a rule of thumb, features of the
 219 impulse response that are *wider* than the event are recognizable in the response
 220 of the filter to the event. Features that are *narrower* (for example the one-sample
 221 impulse at the beginning of the impulse response of the high-pass and notch filters
 222 in Fig. 3) may appear smoothed.

223 **Step.** Certain brain events can be modeled as a step function, for example the
 224 steady-state pedestal that may follow the onset of a stimulus (Picton et al., 1978;
 225 Lammertmann and Lütkenhöner, 2001; Southwell et al., 2017). Figure 7 illus-
 226 trates the various ways a step can be affected by filtering: the step may be *smoothed*
 227 and spread over time, implying that its temporal location is less well defined, and

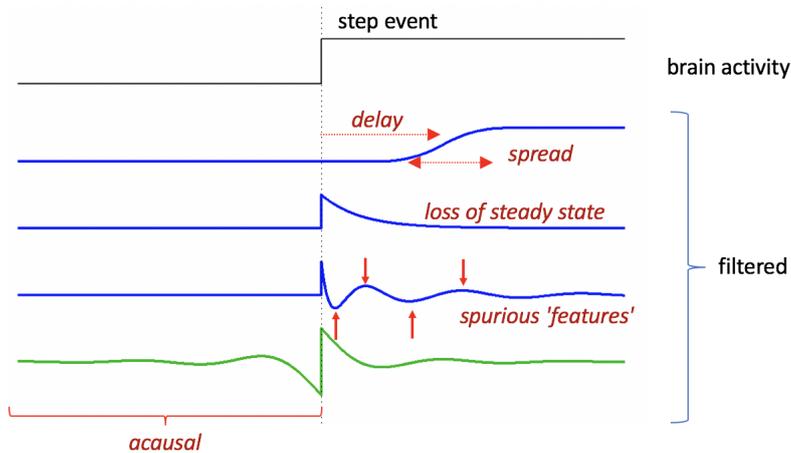


Figure 7: Effects of filtering on a step-like event. The response may be smoothed (low-pass filter), and delayed (causal filter). The steady-state part may be lost (high-pass filter), and spurious 'features' may appear, some of which may occur before the event (acausal filter, green).

228 it may be *delayed* if the filter is causal. Multiple *spurious events* may appear, some
 229 of which may occur before the event if the filter is acausal.

230 The nature of these effects depends on the filter and can be inferred from its
 231 *step response* (integral over time of the impulse response). Step responses of
 232 typical filters are shown in Fig. 8. The sharp transition within the waveform is
 233 smoothed by a low-pass filter (A-B) and delayed relative to the event if the filter is
 234 causal (A), or else it starts before the event if the filter is acausal (B). The steady-
 235 state pedestal is lost for a high-pass (C-E) or bandpass (F-H) filter. The response
 236 may include spurious excursions, some of which precede the event if the filter is
 237 acausal. The response may be markedly oscillatory, in particular for a band-pass
 238 filter (F-H), and it may extend over a remarkably long duration if the filter has a
 239 narrow transfer function.

240 Of course, actual step-like brain events differ from an ideal step. As a rule of
 241 thumb, features of the step response that are *wider* than the event onset will be rec-
 242 ognizable in the output, whereas features that are *narrower* will appear smoothed.
 243 Note that a response of opposite polarity would be triggered by the *offset* of a
 244 pedestal.

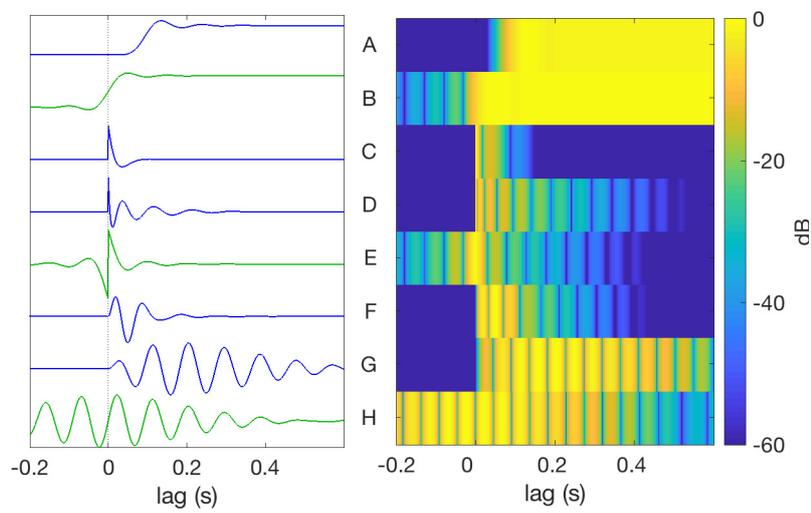


Figure 8: Step responses of typical filters. The left panel shows the step response time series, the right plot shows the log of its absolute value (truncated at -60 dB) on a color scale to reveal the low-amplitude tail. A: Low-pass Butterworth order 8, cutoff 10 Hz, B: same, applied using `filtfilt`, C: High-pass Butterworth order 2, cutoff 10 Hz, D: same, order 8, E: same as D, applied using `filtfilt`, F: Band-pass Butterworth order 2, 10-20 Hz, G: same, 10-12 Hz, H: same as G, applied using `filtfilt`. Filters in green are acausal.

245 **Oscillatory pulse** Some activity within the brain is clearly oscillatory (Buzsáki,
246 2006; da Silva, 2013). The onset or offset of such activity can be modeled as
247 the response to an oscillatory pulse. As Fig. 9 shows, the time course of such a
248 pulse is affected by filtering: it is always *smoothed* and spread over time, it may
249 be *delayed* if the filter is causal, or else start earlier than the event if the filter is
250 acausal. These effects are all the more pronounced as the filter is narrow (as one
251 might one might want to use to increase the SNR of such oscillatory activity).

252 For a notch filter tuned to reject the pulse frequency, ringing artifacts occur
253 at both onset and offset. If the filter is acausal, these artifacts may both precede
254 and follow onset and offset events. For a notch filter tuned to reject power line
255 components (50 or 60 Hz), such effects might also be triggered by fluctuations
256 in amplitude or phase. They might also conceivably affect the shape of a short
257 narrow-band gamma brain response in that frequency region (Fries et al., 2008;
258 Saleem et al., 2017).

259 **5 What can go wrong?**

260 The use of filters raises many concerns, some serious, others merely inconvenient.
261 It is important to understand them, and to report enough details that the reader
262 too fully understands them. An obvious concern is *loss of useful information*
263 suppressed together with the noise. Slightly less obvious is the *distortion* of the
264 temporal features of the target: peaks or transitions may be smoothed, steps may
265 turn into pulses, and *artifactual features* may appear. Most insidious, however, is
266 the *blurring of temporal or causal relations* between features within the signal,
267 or between the signal and external events such as stimuli. This section reviews
268 a gallery of situations in which filtering may give rise to annoying or surprising
269 results.

270 **Loss of information.** This is an obvious gripe: information in frequency ranges
271 rejected by the filter is lost. High-pass filtering may mask slow fluctuations of
272 brain potential, whether spontaneous or stimulus-evoked (Picton et al., 1978; Lam-
273 mertmann and Lütkenhöner, 2001; Vanhatalo et al., 2005; Southwell et al., 2017).
274 Low-pass filtering may mask high-frequency activity (e.g. gamma or high-gamma
275 bands), or useful information about the shape of certain responses (Cole and

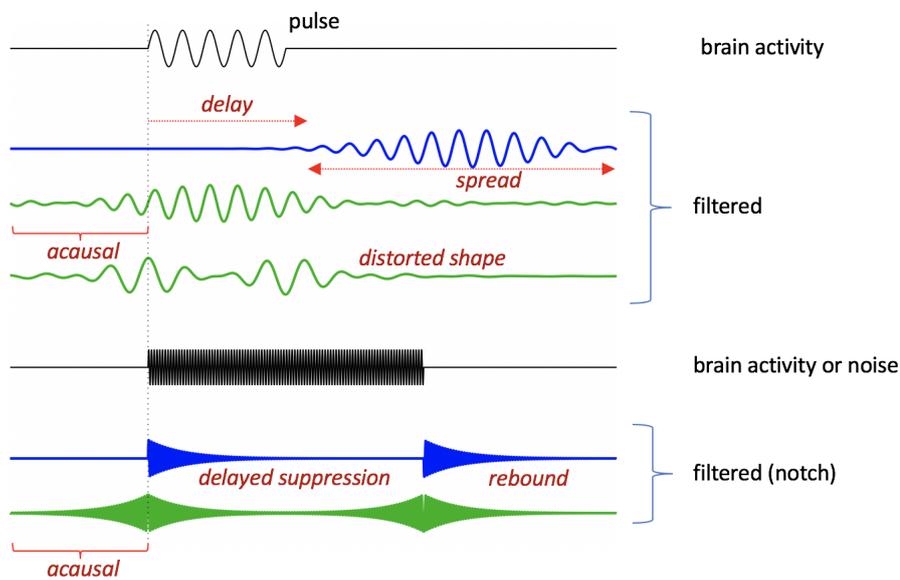


Figure 9: Effects of filtering on a sinusoidal pulse. The pulse is widened and delayed (causal filter) by a bandpass filter. The delay is avoided with a zero-phase filter, but the response then starts before the event. The pattern of distortion may be more complex (here a bandpass with cutoff slightly below the pulse frequency). For a notch filter tuned to the frequency of the pulse, the suppression may be delayed and there may be a rebound artifact after the pulse. If the filter is acausal a 'rebound' artifact may also occur before onset and offset events.

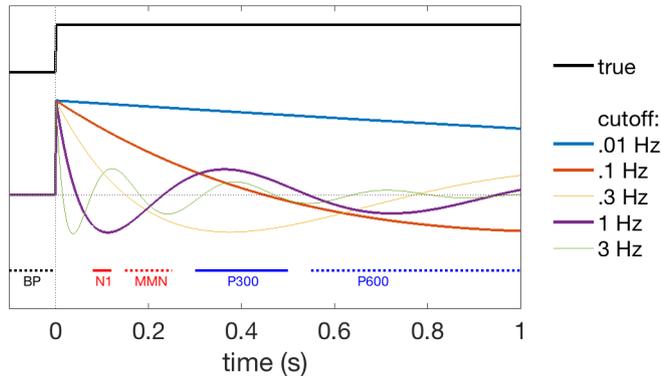


Figure 10: High-pass filter response to a step (Butterworth order 8, cutoff as indicated in legend). The lines at the bottom of each graph indicate temporal intervals within which certain widely-reported ERP features are expected (Berenscheiftspontential BP, N1 or N100, MMN, P300, P600). These were selected for illustrative purposes; numerous other features have been reported in this range.

276 Voytek, 2017; Lozano-Soldevilla, 2018). A notch filter might interfere with nar-
 277 rowband gamma activity that happens to coincide with the notch frequency (Fries
 278 et al., 2008; Saleem et al., 2017). A bandpass filter may reduce the distinction
 279 between shapes of spikes emitted by different neurons and picked up by an extra-
 280 cellular microelectrode, degrading the quality of spike sorting.

281 **Artifactual features** Slightly less obvious is the *distortion* of the temporal fea-
 282 tures of the target: peaks or transitions may be smoothed, steps may turn into
 283 pulses, and so-on. *Artifactual features* may emerge, such as response peaks, or
 284 oscillations (“ringing”) created de novo by the filter in response to some feature
 285 of the target or noise signal. Figure 10 shows the response to a step of a high-pass
 286 filter (Butterworth order 8) of various cutoff frequencies. The response includes
 287 multiple excursions of both polarities (“positivities” and “negativities”) that may
 288 have no obvious counterpart in the brain signal. Disturbingly, the latencies of
 289 some fall in the range of standard ERP response features (schematized as lines in
 290 Figure 10).

291 The morphology of these artifacts depends on both the filter and the brain ac-
 292 tivity, as further illustrated in Fig. 11. An investigator, or a reader, might wrongly
 293 be tempted to assign to the multipolar deflections of the filter response a se-

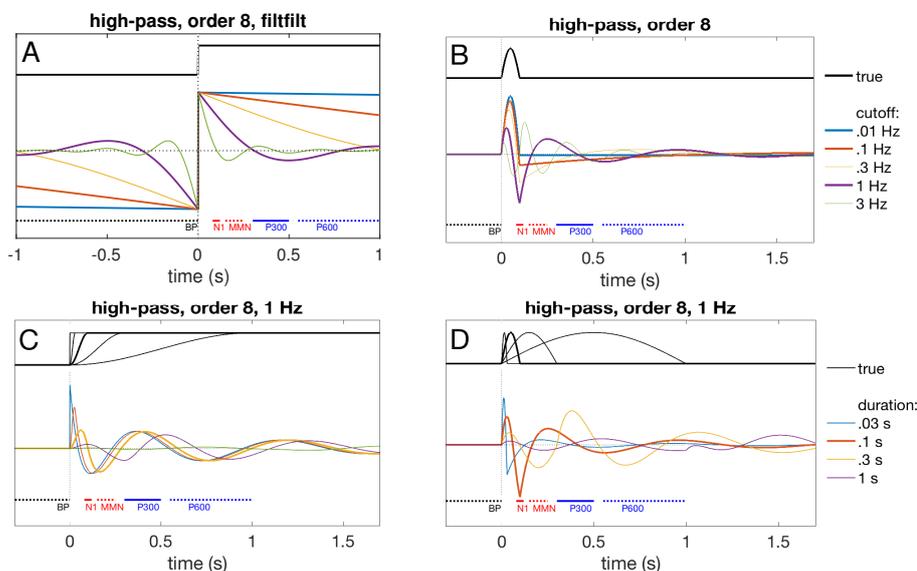


Figure 11: Examples of artifactual features that can arise due to high-pass filtering. A: step response of a zero-phase filter (Butterworth order 8 applied using `filtfilt`), cutoff as indicated in legend. B: response to a pulse of a Butterworth order 8 filter, cutoff as indicated in legend. C: response to a smoothed step of a Butterworth filter of order 8 and cutoff 1 Hz, step transition duration as indicated in legend. D: response to a pulse of a Butterworth filter of order 8 and cutoff 1 Hz, pulse duration as indicated in legend. The lines at the bottom of each graph are as defined in Fig. 10.

294 quence of distinct physiological processes. Similar issues have been pointed out
 295 with respect to spike waveform morphology from extracellular recordings (Quian
 296 Quiroga, 2009; Molden et al., 2013).

297 **Spurious oscillations** Oscillatory phenomena play an important role in the brain
 298 (Buzsáki, 2006; da Silva, 2013), and many response patterns are interpreted as
 299 reflecting oscillatory activity (Zoefel and VanRullen, 2017; Meyer, 2017; Singer,
 300 2018), although in some cases this interpretation has been questioned (Yeung et
 301 al., 2004; Yuval-Greenberg et al., 2008; Jones, 2016; van Ede et al., 2018).

302 Non-oscillatory inputs (e.g. an impulse or step) can trigger a filter response
 303 with distinctly oscillatory features. Figure 12 shows the response of a 8-11 Hz
 304 bandpass filter (such as might be used to enhance alpha activity relative to back-
 305 ground noise) to several inputs, including a 10 Hz sinusoidal pulse (top) and two

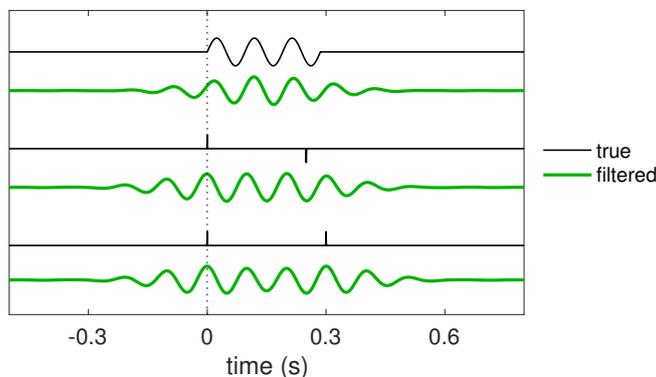


Figure 12: Response of a bandpass filter (8-11 Hz, Butterworth order 8 applied with `filtfilt`) to a 10 Hz sinusoidal pulse (top) and to an input consisting of two impulses (middle and bottom).

306 configurations of impulses. Visually, the responses to the non-oscillatory impulse
 307 pairs are, if anything, more convincingly oscillatory than the response to the os-
 308 cillatory input!

309 Oscillations tend to occur with a frequency close to a filter cutoff, and to be
 310 more salient for filters with a high order (and steeper cutoffs of the transfer func-
 311 tion). They can occur for any filter with a sharp cutoff in the frequency domain,
 312 and are particularly salient for band-pass filters, as high-pass and low-pass cutoffs
 313 are close and may interact. Furthermore, if the pass band is narrow, the inves-
 314 tigator might be tempted to choose a filter with steep cutoffs, resulting in a long
 315 impulse response.

316 **Masking or reintroduction of artifacts** Cognitive neuroscientists are alert to
 317 potential artifacts, for example muscular activity that differs between conditions
 318 due to different levels of effort. Muscular artifacts are most prominent in the
 319 gamma range (where they emerge from the $1/f$ background), and thus low-pass
 320 filtering is often indicated to eliminate them. Indeed, visually, there is little in
 321 the low-pass filtered signal to suggest muscle artifacts. Low-frequency correlates
 322 are nonetheless present (muscle spikes are wideband) and could potentially in-
 323 duce a statistically significant difference between conditions. Filtering masks this
 324 problem (if there is one).

325 Conditions that require different levels of effort might also differ in the num-
 326 ber of eye-blinks that they induce. Subjects are often encouraged to blink between

327 trials, so as avoid contaminating data within the trials. However if high-pass or
328 band-pass filtering is applied to the data before cutting them into epochs, the fil-
329 ter response to the blink may extend into the epoch, again inducing a statistically
330 significant difference between conditions. For a causal filter each epoch is con-
331 taminated by any blinks that precede it, for an acausal filter it may also be con-
332 taminated by any blinks that follow it.

333 **Temporal blurring, delay, causality** The most subtle effect of filtering is the
334 blurring of temporal relationships, which can interfere with the comparison be-
335 tween brain measurements and stimulation or behavior, or between recordings
336 at different recording sites, or between different frequency bands. Temporal or
337 causal relationships between events are less clear when looking at filtered data.
338 The problem is mild if the filter impulse response is short relative to the phe-
339 nomena being measured, but such is not always the case. Impulse responses of
340 commonly-used filters may extend over hundreds of milliseconds (Fig. 6) whereas
341 important stages of neural processing may occur over 1-10 milliseconds (e.g.
342 Schmolesky et al., 1998).

343 The time course of sensory processing is often inferred either from the latency
344 of the *peak* response to stimulation, or of the point at which the response emerges
345 from background noise. A causal filter introduces a systematic bias in the first
346 measure (towards a longer latency), and an acausal filter a bias in the second
347 measure (towards a shorter latency). The early part of an acausal filter response
348 might misleadingly masquerade as an early brain response, or as the correlate of
349 a predictive mechanism.

350 Similar issues arise for Temporal Response Functions (TRF) obtained by fit-
351 ting stimulus and response data with a linear model (Lalor and Foxe, 2010; Ding
352 and Simon, 2013; O’Sullivan et al., 2015; Crosse et al., 2016). TRF analysis has
353 become popular as a tool to characterize the response to continuous stimuli such
354 as speech or environmental sound. Features of the TRF (e.g. peaks) are some-
355 times interpreted as reflecting particular brain processes, and inferences are made
356 about their anatomical localization based on their latency. If, as is common, the
357 brain data are filtered to restrict the analysis to a frequency range where the re-
358 sponse is expected to best follow the stimulus (e.g. 1-10 Hz), the estimated TRF
359 will approximate the *convolution of the real TRF with the impulse response of the*

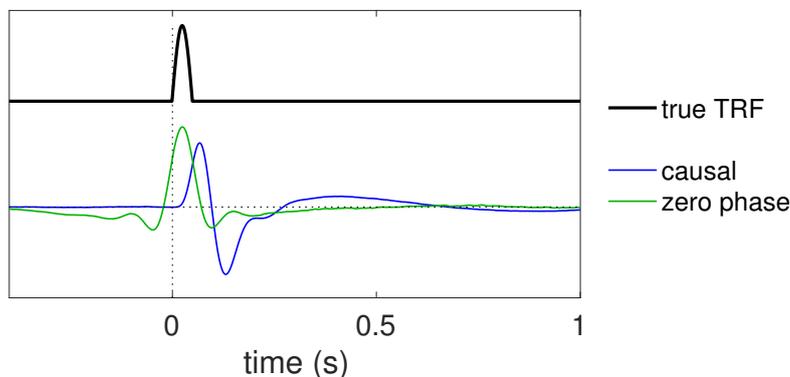


Figure 13: Temporal Response Function estimated from simulated stimulus-response data. Black: “true” TRF. Thick blue: TRF estimated using response data that has been filtered by a causal filter (Butterworth bandpass 1-10 Hz, order 4+4). Green: same with acausal filter (Matlab’s `filtfilt`).

360 *filter*. To illustrate this point, a simulated “stimulus” consisting of Gaussian white
 361 noise was processed with a simulated “TRF” consisting of a half-sinusoidal pulse
 362 of duration 50 ms (Fig. 13 black line) to obtain simulated brain data. These data
 363 were then filtered with a bandpass filter and the TRF estimated using the mTRF
 364 toolbox (Crosse et al., 2016). Figure 13 (blue line) shows the TRF estimate. The
 365 green line is the estimate when the same filter was applied in both directions using
 366 `filtfilt`. In both cases, the shape of the estimated TRF differs from that of the
 367 real TRF. The potential effect of filtering on TRFs is rarely discussed, and filters
 368 used to preprocess the data are often not fully described.

369 **Time-frequency analysis** Time-Frequency (TF) analysis is usually seen as a
 370 data analysis rather than filtering tool, nonetheless filters are involved “under the
 371 hood” and TF representations are vulnerable to similar problems as noted for fil-
 372 ters.

373 Time-frequency (TF) representations (e.g. spectrograms) are obtained by ap-
 374 plying short-term spectral analysis to the data with a short analysis window that
 375 slides in time. At each time point the analysis yields a spectrum and these spectra
 376 are concatenated to form the two-dimensional TF representation. Each pixel in
 377 the 2-D representation is indexed by time (abscissa) and analysis frequency (ordi-
 378 nate). The representation usually displays some transform of amplitude or power
 379 (Fig. 14 B, C), but it is also possible to plot phase (Fig. 14 D, E).

380 In a standard “short-term Fourier transform” (STFT) spectrogram, the size of
381 the analysis window is the same for all frequencies (Fig. 14 B). In contrast, in a
382 wavelet spectrogram this parameter varies with frequency, for example such that
383 each analysis window spans the same number of cycles (Fig. 14 C), i.e. it is longer
384 at low frequencies and shorter at high frequencies.

385 The value of the TF representation at the analysis time point reflects all sig-
386 nal values within the analysis window. Conversely each signal value impacts TF
387 values over a range of analysis time points. The overall alignment between data
388 values and TF values depends on the convention chosen to assign a time index to
389 the analysis value. TF samples can be aligned with the *end* of the analysis win-
390 dow, corresponding to a causal analysis, or more commonly with the *center* of the
391 analysis window, corresponding to an acausal analysis. TF features are thus ei-
392 ther delayed relative to events within the data (causal analysis) or else they partly
393 reflect future events (acausal analysis).

394 Figures 14 (B, C) show TF magnitude representations in response to a pulse-
395 shaped input signal. The temporally-localized event at $t = 0$ affects the spec-
396 trogram over a range of time points spanning the event (for example ± 0.25 s in
397 Fig. 14 B). Equivalently, the value of the spectrogram at $t = 0$ can “see” all signal
398 events within a range of time points spanning that instant.

399 Figures 14 (D, E) show TF phase representations in response to the same
400 pulse-shaped input signal. Phase is defined only for non-zero magnitude, i.e. only
401 when the pulse falls within the analysis window. The event at $t = 0$ affects the
402 phase estimate of analyses made over a range of time points spanning the event.
403 Equivalently, the phase estimate obtained at $t = 0$ is affected by all events within
404 that range, some of which occur *later* than the analysis point. This blurred, non-
405 causal relation between data and TF analysis can lead to misleading conclusions.

406 As an example of such a misleading conclusion, suppose that we wish to es-
407 tablish whether the phase of brain oscillations preceding a stimulus predicts the
408 brain or behavioral response to that stimulus. TF analysis seems to be the right
409 tool for that purpose. Indeed, using it we observe that phase within some fre-
410 quency band (e.g. alpha) measured just before the trial is systematically biased
411 towards a particular value on successful trials. From this we conclude that oscil-
412 latory phase preceding stimulation determines the response. Unfortunately, that
413 conclusion is not warranted if the analysis window overlaps the stimulus-evoked

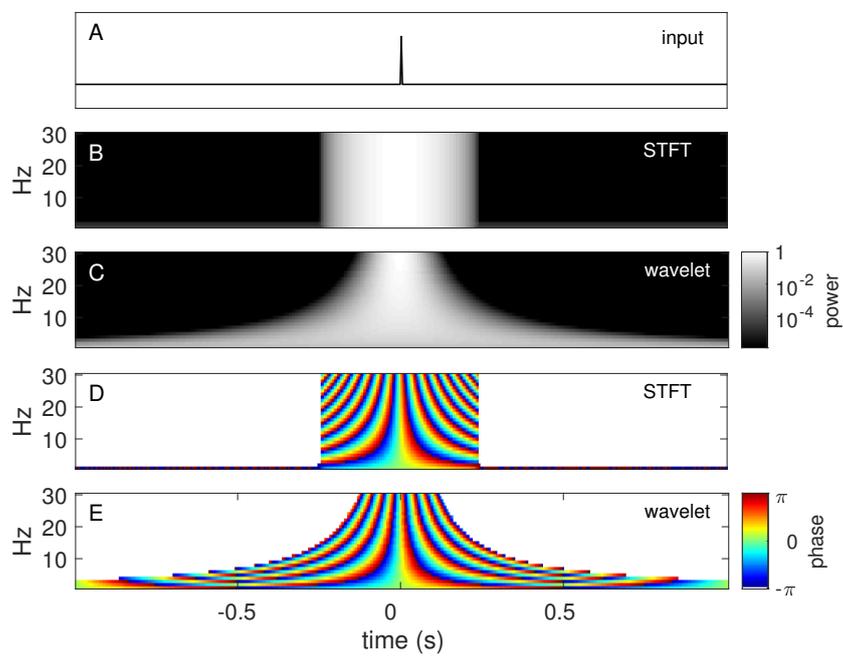


Figure 14: Time frequency power spectral density plots in response to a Dirac pulse (A). B: STFT-based spectrogram with an analysis window of duration 0.5 s. C: wavelet-based spectrogram with an analysis window of duration 7 cycles. D and E: phase plots corresponding to B and C.

414 sensory or behavioral motor response. The interesting conclusion (response de-
415 pendency on prior phase) can only be made if that more trivial possibility is ruled
416 out, for example by using causal TF analysis (Zoefel and Heil, 2012). Similar is-
417 sues may arise in analyses of cross-spectral coupling. These issues may be harder
418 to spot if wavelet analysis is involved.

419 **Notch filter artifacts** A narrow notch filter works well to remove narrowband
420 interference that is stationary (e.g. 50 or 60 Hz line power). Additional notches
421 may be placed at harmonics if needed. However, notch filtering may be less ef-
422 fective if the interference is not stationary. Amplitude fluctuations may occur if
423 the subject moves, and for MEG the phase may fluctuate with changes in load
424 in the tri-phase power network from which originates the interference. Artifacts
425 can also be triggered by large-amplitude glitches (Kıraç et al., 2015). Filtering is
426 ineffective in removing interference close to the ends of the data (Fig. 9, bottom),
427 and should not be applied to epoched data.

428 **Inadequate antialiasing.** Effects of the antialiasing filter are rarely noticed or
429 objectionable. More serious may be a *lack* of sufficient antialiasing. Figure 15
430 (left) shows the power spectrum of a sample of data from a MEG system with
431 shallow (or missing) antialiasing. As common in MEG, there are salient power
432 line components at 60 Hz and harmonics (black arrows), but also many additional
433 narrowband components that likely reflect aliasing of sources with frequencies be-
434 yond the Nyquist frequency (250 Hz). Possible sources include higher harmonics
435 of 60 Hz, or high-frequency interference from computer screens, switching power
436 supplies, etc. The frequency of the artifact cannot be known for sure. For exam-
437 ple, the spectral line at 200 Hz (red arrow) could be the aliased 300 Hz harmonic
438 of the power line interference, or it could have some other origin. This example
439 underlines the importance of an adequate antialiasing filter.

440 In contrast, Fig. 15 (right) shows the response to a sharp change in sensor state
441 of one channel of a different MEG system with a particularly steep antialiasing
442 filter (8th order elliptic filter with 120 dB rejection and 0.1 ripple in pass-band)
443 (Oswal et al., 2016). The data show a prominent oscillatory pattern that is likely
444 not present in the magnetic field measured by the device.

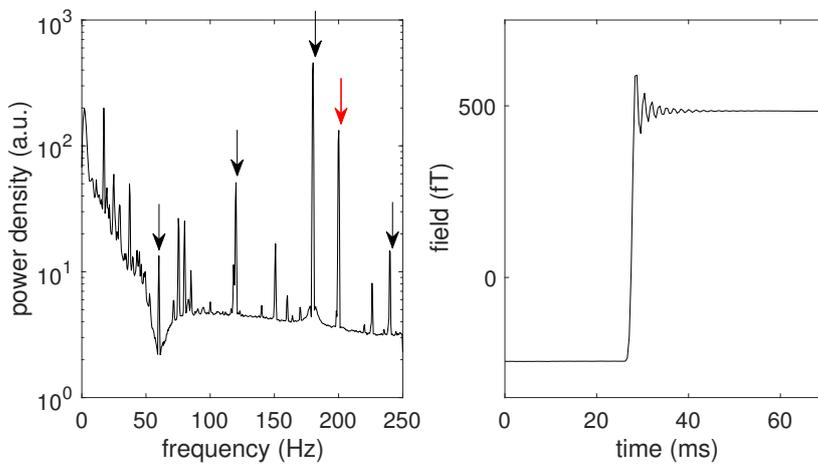


Figure 15: *Left: power spectrum of data from a MEG system with a shallow (or missing) antialiasing filter. Peaks at 60, 120, 180 and 240 Hz (black arrows) probably reflect power line harmonics, but the origin of the other peaks is mysterious. They might be the result of aliasing of high-frequency sources within the environment, or of higher harmonics of 60 Hz. For example the peak at 200 Hz (red arrow) might result from aliasing of the fifth harmonic of 60 Hz. Right: ringing artifact in a MEG system with a particularly steep antialiasing filter. The magnetic field change was a sharp step.*

445 **6 How to fix it ?**

446 **Report full filter specs** This should go without saying. In each case, the prob-
447 lem is compounded if the reader can't form an opinion about possible effects of
448 filtering. Filter type, order, frequency parameters, and whether it was applied
449 in one direction or both (`filtfilt`) should be reported. Include a plot of the
450 impulse response (and/or step response) as an inset to one of the plots.

451 **Antialiasing artifacts.** Antialiasing artifacts are rarely an issue. In the event
452 that they are, consider first whether antialiasing is needed. If you are certain
453 that the original data contain no power beyond the Nyquist frequency, omit the
454 filter and live dangerously. If instead there are high-frequency sources of large
455 amplitude, you might want to verify that the antialiasing filter attenuates them
456 sufficiently before sampling. Note that, because of the aliasing, the frequency
457 of those sources cannot be inferred with confidence. A wide-band oscilloscope
458 or spectrum analyzer might be of use. To reduce temporal smear and ringing,
459 consider using an antialiasing filter with a lower cutoff and shallower slope. In the
460 case of downsampling or resampling, consider alternatives such as interpolation
461 (e.g. linear, cubic or spline). The artifact of Fig. 15 (right) can be removed as
462 described by de Cheveigné and Arzounian (2018).

463 **Low-pass** First, ask whether the aim is to *smooth* the temporal waveform, for
464 example to enhance the clarity of a plot, or whether it is to ensure attenuation of
465 high-frequency power (for example preceding downsampling). If the former, con-
466 sider using a simple smoothing kernel, for example square, triangular, or Gaussian
467 (Fig. 6 a, b). Such kernels have a limited and well-defined temporal extent, and no
468 negative portions so they do not produce ringing. They tend however to have poor
469 spectral properties. Conversely, if temporal distortion is of no importance, the fil-
470 ter can be optimized based only on its frequency response properties (Widmann
471 et al., 2015).

472 If data are recorded on multiple channels (e.g. local field potentials, EEG,
473 or MEG), spatial filters may be applied to remove noise sources with a spatial
474 signature different from the target sources. The appropriate filters can be found
475 based on prior knowledge or using data-driven algorithms (e.g. Parra et al., 2005;

476 de Cheveigné, 2016).

477 **High-pass** If the high-pass filter is required merely to remove a constant DC
478 offset, consider subtracting the overall mean instead. If there is also a slow trend,
479 consider *detrending* rather than high-pass filtering. Detrending involves fitting
480 a function (slowly-varying so as to fit the trend but not faster patterns) to the
481 data and then subtracting the fit. Suitable functions include low-order polynomi-
482 als. Like filtering, detrending is sensitive to temporally localized events such as
483 glitches, however these can be addressed by *robust detrending* (de Cheveigné and
484 Arzounian, 2018).

485 If the slow trend signal can be estimated independently from the measurement
486 that it contaminates, consider using regression techniques to factor it out (Vrba
487 and Robinson, 2001; de Cheveigné and Simon, 2007). Even when this is impossi-
488 ble, if the data are multichannel, consider using a component-analysis technique
489 to factor it out, as has also been suggested to obtain distortion-free extracellular
490 spike waveforms (Molden et al., 2013).

491 If all else fails, and high-pass filtering must be used, pay particular attention
492 to its possible effects on the morphology of responses. If the initial portion of the
493 data (duration on the order of $1/f_c$ where f_c is the cutoff frequency) is on average
494 far from zero, it may be useful to subtract the average over that portion, so as to
495 minimize the filter response to the implicit initial step (the filter treats the input
496 data as being preceded by zeros). If the data are to be cut into epochs (e.g. to
497 excise responses to repeated stimuli), it is usually best to filter the continuous data
498 first. Be aware that artifacts from out-of-epoch events (e.g. eye blinks) may extend
499 to within the epoch.

500 **Band-pass** Consider whether a band-pass filter is really needed, as the potential
501 for artifactual patterns is great. If band-pass filtering must be applied (for ex-
502 ample to improve signal-to-noise ratio to assist a component-analysis technique),
503 consider filters with relatively shallow slopes, and cutoff frequencies distant from
504 the activity of interest. Be on the lookout for artifactual results due to the filtering.

505 **Notch** Notch filtering is usually motivated by the desire to suppress line noise
506 (50 Hz or 60 Hz and harmonics). Of course, the best approach is to eliminate that

507 noise at the source by careful design of the setup, but this is not always feasible.
508 As an alternative to filtering, it may be possible to measure the line noise on one
509 or more reference channels and regress them out of the data (Vrba and Robin-
510 son, 2001). If the data are multichannel, consider using component analysis to
511 isolate the line noise components and regress them out (Delorme et al., 2012; de
512 Cheveigné and Parra, 2014; de Cheveigné and Arzounian, 2015).

513 If the high-frequency region is not of interest, a simple expedient is to apply
514 a boxcar smoothing kernel of size 1/50 Hz (or 1/60 Hz as appropriate). This
515 simple low-pass filter has zeros at the line frequency and all its harmonics, and
516 thus perfectly cancels line noise. The mild loss of temporal resolution (on the
517 order of 20 ms) might be deemed acceptable. If the sampling rate differs from a
518 multiple of the line frequency, the appropriate kernel can be implemented using
519 interpolation (see de Cheveigné and Arzounian, 2018, for details).

520 **Time-frequency analysis** If the patterns of interest can be interpreted in the
521 time domain, eschew TF analysis. If the data are multichannel, and the aim is
522 to increase the signal-to-noise ratio of narrow-band or stimulus-induced activity,
523 consider component analysis techniques that can boost SNR of narrow-band sig-
524 nals (Nikulin et al., 2011; de Cheveigné and Arzounian, 2015).

525 If TF analysis must be applied, consider using fixed kernel-size analysis (e.g.
526 DFT) rather than, or in addition to, wavelet analysis, so that temporal bias and
527 smearing are uniform across the frequency axis. Consider using relatively short
528 analysis windows to reduce temporal bias and/or smearing. Weigh carefully the
529 choice between causal analysis (temporal bias but no causality issues) and acausal
530 analysis (no temporal bias but risk of misleading causal relations). In every case,
531 be alert for potential artifacts. One should be particularly concerned if an interest-
532 ing effect only emerges with a particular analysis method.

533 **7 Horror scenarios**

534 This section imagines scenarios in which filtering effects might affect the science.
535 Some are mildly embarrassing, others might keep a scientist awake at night.

536 **Missed observation.** Researcher **A** applies a high-pass filter to data recorded
537 over a long period and fails to notice the existence of infra-slow brain activity (as
538 reported by Vanhatalo et al., 2005). Researcher **B** applies a low-pass filter and
539 fails to notice that a certain oscillatory activity is not sinusoidal (as reported by
540 Cole and Voytek, 2017). It is frustrating to miss part of the phenomena one set out
541 to study.

542 **Bias from eye movements.** Following a scenario hinted at in Sect. 5, researcher
543 **C** runs a study in which some conditions are more demanding than others. Sub-
544 jects are instructed to blink only between trials, but because acausal high-pass (or
545 band-pass) filtering is applied to the data, each blink triggers a filter response that
546 extends into the trial, resulting in a significant difference between conditions. Re-
547 searcher **D** runs studies that create miniature eye movements (microsaccades) that
548 differ between conditions. Microsaccades introduce so-called spike potentials,
549 transients with a time course of a few tens of milliseconds, which after TF anal-
550 ysis boost energy in the gamma band selectively in some conditions rather than
551 others (Yuval-Greenberg et al., 2008). In both cases ocular activity masquerades
552 as brain activity.

553 **Distorted observation.** Researcher **E** records brain responses to stimulation,
554 applies a high-pass filter to attenuate a pesky slow drift, and fails to notice that
555 the brain response actually consisted of a sustained pedestal. Instead, a series of
556 positive and negative peaks is observed and interpreted as reflecting a succession
557 of processing stages in the brain. In a milder version of this scenario, the brain
558 response does include such peaks, but the filter affects their position, leading to
559 incorrect inferences concerning brain processing latencies.

560 **Flawed replication.** Researcher **F** replicates Researcher **E**'s experiments, using
561 the same filters and generating the same artifacts. Results are consistent, giving
562 weight to the conclusion that they are real.

563 **Faulty communication.** Researcher **G**, who is filter-savvy, reads his/her col-
564 league's papers and suspects something is amiss, but cannot draw firm conclu-
565 sions because methods were not described in full. He/she re-runs the experiments

566 with careful methods, and finds results that invalidate the previous studies. The
567 paper is not published because the study does not offer new results.

568 **Proliferation of “new” results.** Other researchers run further studies using anal-
569 ogous stimuli, but using different analysis parameters. New patterns of results are
570 found that are interpreted as new discoveries, whereas the actual brain response
571 (in this hypothetical scenario) is the same.

572 **Oscillations?** Researcher **H** knows that with the right kind of preprocessing,
573 multiple layers of oscillatory activity can be found hidden within brain signals,
574 and believes that the analysis is revealing them. Researcher **I** suspects that these
575 oscillations reflect filter ringing, but finds it hard to counter **H**’s arguments (Fourier’s
576 theorem says that the oscillations are indeed there). **I** remains worried because the
577 observed oscillations depend on the choice of filter, but **H** is not: different filters
578 extract different parts of the data, each with its own oscillatory nature. The debate
579 mobilizes a good proportion of their energy.

580 **Biased time-frequency analysis.** Researcher **K** uses time-frequency analysis to
581 test the hypothesis that the phase of ongoing brain oscillations modulates percep-
582 tual sensitivity. To avoid contamination by the sensory or behavioral response,
583 the analysis is carefully restricted to the data preceding stimulation. However the
584 analysis window, centered on the analysis point, extends far enough to include
585 the sensory or behavioral response, biasing the distribution of measured phase. **K**
586 concludes (incorrectly in this hypothetical scenario) that the hypothesis is correct.
587 In a variant of this scenario, **L** uses time frequency analysis to test the hypothe-
588 sis that brain activity is durably entrained by a rhythmical stimulus. The analysis
589 is applied to the data beyond the stimulus offset, but the analysis window over-
590 laps with the stimulus-evoked response, again biasing the phase distrbution. **L**
591 concludes (again incorrectly in this scenario) that the hypothesis was correct.

592 **8 Discussion**

593 A filter has one purpose, improve SNR, and two effects: improve SNR and distort
594 the signal. Many investigators consider only the first and neglect the second. The

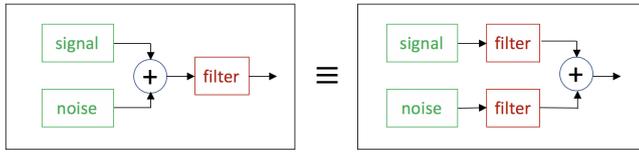


Figure 16: Linear operations can be swapped. The filtered noisy signal is the superposition of the filtered signal and the filtered noise.

595 filtered data are the sum of the filtered target signal and the filtered noise, and thus
 596 one can focus separately on these two effects (Fig. 16). Here we focused on target
 597 distortion.

598 Issues related to distortion have been raised before, in particular distortion due
 599 to low-pass filtering (VanRullen, 2011; Rousselet, 2012; Widmann and Schröger,
 600 2012), high-pass filtering (Kappenman and Luck, 2010; Acunzo et al., 2012; Tan-
 601 ner et al., 2015, 2016; Widmann et al., 2015; Lopez-Calderon and Luck, 2014)
 602 and band-pass filtering (Yeung et al., 2004), in the context of EEG and MEG and
 603 also extracellular recordings (Quiñero Quiroga, 2009; Molden et al., 2013; Yael and
 604 Bar-Gad, 2017). They are also discussed in textbooks and guidelines (Picton et
 605 al., 2000; Nunes and Srinivasan, 2006; Gross et al., 2013; Keil et al., 2014; Luck,
 606 2014; Puce and Hämäläinen, 2017; Cohen, 2014, 2017).

607 **How serious are these issues?** They can be recapitulated as follows. First, the
 608 *loss of information* in spectral regions suppressed by the filter. This problem is
 609 straightforward and does not need elaboration. Second, the *distortion of response*
 610 *waveforms* and the *emergence of spurious features*. This is certainly a concern if
 611 spurious features (e.g. delayed excursions, or ringing) misleadingly suggest brain
 612 activity that is not there. Third, the *blurring of temporal relations*, in particular
 613 violation of *causality*. This too is a concern given the importance of response
 614 latency in inferring the sequence of neural events, or the anatomical stage at which
 615 they occur. Fourth, *the non-uniqueness of phenomenological descriptions*: the
 616 same event can take very diverse shapes depending on the analysis. This can
 617 interfere with comparisons between studies, and can lead to redundant reports of
 618 the same phenomenon under different guise. Fifth, the *lack of details* required
 619 by a knowledgeable reader to infer the processing involved. Rather than an issue
 620 with filtering per se, the issue is with sloppy practice in reporting methodological

621 details of filtering and TF analysis.

622 Cutoff frequencies may be reported, but not the type of filter, its order, or
623 whether it was applied in a single pass or both ways. As illustrated in Figs. 6
624 and 8, the cutoff frequency of a filter is not sufficient to characterize its impulse
625 or step response, information that is needed to guess how it might have impacted
626 a reported response. Failure to report details can be due to space limits (some-
627 times misguidedly imposed by journals), incomplete knowledge (e.g. proprietary
628 or poorly documented software), reluctance to appear pedantic by reporting mun-
629 dane trivia, or lack of understanding that this information is important.

630 The issue of non-uniqueness is not often raised. Non-uniqueness refers to the
631 fact that analysis of the same phenomenon can give rise to different descriptions
632 depending on the analysis parameters, making it hard to compare across studies. It
633 is sometimes recommended that parameters should be adjusted to the task at hand,
634 rather than use default values proposed by the software (Widmann and Schröger,
635 2012). Optimizing data analysis is laudable, but it carries the risks of “cherry-
636 picking” or “double-dipping” (Kriegeskorte et al., 2009).

637 **Quid frequency and phase?** Filter design has developed sophisticated methods
638 to optimize the frequency response to maximize rejection, minimize ripple, and/or
639 obtain the steepest possible transition between pass and stop bands. Engineers and
640 scientists trained in those methods tend to choose a filter based on these properties,
641 with less attention to their time-domain counterpart. It is not always clear that this
642 emphasis is justified. For example, a band-pass filter with steep slopes might be
643 motivated by the desire to “keep the delta band distinct from the theta band”,
644 but given that there is little theoretical or phenomenological evidence for a clear
645 boundary between bands, this is should perhaps not be a primary goal.

646 For any given magnitude response, there are multiple filters with different
647 phase responses. Of particular interest are *zero phase* filters with minimal wave-
648 form distortion and no delay (but that are unfortunately acausal), and *minimum*
649 *phase* filters with greater waveform distortion but that are causal. The choice
650 between these phase characteristics (or others) depends on whether one wishes
651 to favour causality, overall delay, or waveform distortion, knowing that it is im-
652 possible to favour all. Some authors recommend causal filters (Rousselet, 2012),
653 others linear phase or acausal (Widmann and Schröger, 2012). Some studies re-

654 port using simple filters (e.g. low order Butterworth), others sophisticated designs
655 (e.g. Chebyshev or elliptic) or even “brickwall” filters implemented in the Fourier
656 domain.

657 A crucial point that we strive to make in this paper is that *no* choice of filter
658 can avoid temporal distortion, as *any* filter entails scrambling of the temporal axis
659 (Fig. 2). Given that a filter with steep slopes in the frequency domain entails a
660 long impulse response, it may be worth relaxing spectral criteria so as to optimize
661 temporal properties.

662 **Causality, again.** As mentioned earlier, for an acausal filter the output depends
663 on input values that occur later in the future. No physical system can have this
664 behaviour. Offline analysis allows us greater flexibility to align the analysis ar-
665 bitrarily with respect to the data, but we must be clear about what this implies.
666 If we wish to relate the “brain response” to other events within the brain or the
667 world (e.g. stimuli or behaviour), acausal filtering implies that that this response
668 might depend on signal samples that occur *after* those events, indeed, a violation
669 of causality.

670 **9 Recommendations**

671 **Document.** This should go without saying, but many (most?) papers provide
672 incomplete information about the filters employed, a situation exacerbated by the
673 insistence of some journals on limiting the space devoted to methods. Data analy-
674 sis decisions can be justified, incomplete reporting cannot. The reader needs this
675 information to infer the brain signal from the patterns reported.

676 To authors: provide full specifications of the filters applied to the data. A
677 simple plot of the impulse response (or step response) as an insert can be very
678 helpful. To editors and reviewers: demand this information. To journals: avoid
679 requirements that discourage proper documentation. To equipment manufactur-
680 ers: provide full specifications of any hardware filters.

681 **Know your filters.** Make sure that you know the exact filters that are involved in
682 your data recording and analysis. This may require delving into the documentation
683 (or even the code) of your analysis software (e.g. EEGLab, FieldTrip, SPM. etc.).

684 Plot the impulse response and/or the step response and paste it on the wall in front
685 of your desk. If several filters are cascaded, plot the response of their cascade. If
686 specs are lacking, figure out how to deliver a pulse (and/or step) to the recording
687 device and plot the resulting response. If you are using TF analysis, do you know
688 exactly what kernels were employed? Are they causal and thus likely to introduce
689 latency? Are they instead acausal (e.g. zero-phase) and thus likely to confuse
690 causal relations? Are they wavelets, in which case temporal spread and latency
691 might differ across frequency bands? All this should be known.

692 **Know your noise.** The main purpose of a filter is to attenuate noise. What is
693 that noise, where does it come from? Might it be possible to mitigate it at the
694 source? Some experimenters speak of their rig as if it were inhabited by gremlins.
695 This deserves little patience: how can one understand the brain if we can't find
696 the source of line noise in the rig? It may not be possible to suppress the noise
697 (e.g. turn off myogenic, cardiac, ocular or alpha activity, tramways in the street,
698 etc.) but at least the source should be understood. Given that signal and noise
699 both impact the results, understanding a noise process merits as much effort as
700 understanding a brain process.

701 **Eliminate noise at the source.** No need for a filter if there is nothing to at-
702 tenuate. To get rid of line noise: banish power cables from the vicinity of the
703 setup, use lights fed with filtered DC, apply proper shielding (electrostatic cou-
704 pling), avoid loops (magnetic coupling), avoid ground loops (ensure that ground
705 cables and shields have a star topology with no loops), etc. To eliminate high-
706 frequency noise: banish computer screens, fluorescent lights, equipment with
707 switching power supplies, cell phones, etc. If need, apply Faraday shielding. To
708 minimize slow drifts in EEG: follow appropriate procedures when applying the
709 electrodes, and keep the subjects cool. To minimize alpha components: ensure
710 that subjects keep their eyes open, give them a task to keep them alert, and so
711 on. Textbooks(e.g. Luck, 2014; Cohen, 2014) and guidelines can offer many such
712 suggestions.

713 **Ensure that you have adequate antialiasing.** Antialiasing filters in recording
714 equipment are not always well documented. In some situations they might prove

715 insufficient if there is high amplitude noise with a frequency beyond the Nyquist
716 rate (for example from a computer screen, fluorescent light, or cell phone). A
717 similar issue may arise when downsampling digital data: does the low-pass filter
718 suffice to ensure that aliased components are negligible? This may require check-
719 ing the data and/or software at hand (at the time of writing, Matlab's `resample`
720 sets the low-pass cutoff *at* Nyquist rather than below, which is inadequate).

721 **Consider alternatives to filtering.** Consider *detrending* (in particular robust de-
722 trending) as an alternative to high-pass filtering (Bigdely-Shamlo et al., 2015; de
723 Cheveigné and Arzounian, 2018). Consider using an independent reference signal
724 measurement that picks up only noise, and use *regression* techniques to factor out
725 the noise (Vrba and Robinson, 2001; de Cheveigné and Simon, 2007; Molden et
726 al., 2013). Consider component analysis techniques to design a *spatial filter* that
727 factors out the noise (Parra et al., 2005; Delorme et al., 2012; de Cheveigné and
728 Parra, 2014).

729 **Choose the right filter.** If filter we must, a prime consideration is whether to
730 optimize the time domain (minimal distortion of the waveform) or the frequency
731 domain (optimal frequency response), the two being at loggerheads. Taking the
732 example of a low-pass filter, if our goal is to smooth the waveform to enhance
733 the visual clarity of a plot, or locate a peak with less jitter, then a simple box-car
734 smoothing kernel (rectangular impulse response) may be sufficient, with minimal
735 temporal blurring. The poor frequency response of such a low-pass filter is of
736 little import. If instead the focus is on spectral features (e.g. frequency-following
737 response, or narrowband oscillations), we may wish to optimize the spectral prop-
738 erties of the filter at the expense of greater temporal smearing. If the focus is on
739 spectrotemporal features, then the choice of filter(s) necessarily involves a tradeoff
740 between the two (Cohen, 2014).

741 **Simulate.** It is hard to fully predict the impact of filtering, particularly if mul-
742 tiple stages are cascaded. A simple expedient is to simulate the situation using
743 a known target signal (e.g. an idealized evoked response) and known noise (e.g.
744 EEG data from an unrelated recording). The effect of filtering can then be evalu-
745 ated separately on each, given that the filtered sum is the sum of the filtered parts

746 (Fig. 16).

747 The synthetic target signal could be an impulse or step (to visualize canonical
748 response properties), or a signal similar to a typically-observed response (to see
749 how processing might affect it), or a signal constructed to mimic the observed
750 response after filtering (to help infer true patterns from observations). Observing
751 the response to the target tells us how it is distorted, observing the response to
752 the noise tells us how well it is attenuated and what artifactual patterns to expect.
753 Comparing the two tells us whether our observation is helped (or hindered) by
754 filtering.

755 **Be paranoid.** Is the effect of interest only visible for a particular type of filter,
756 or a particular variety of TF analysis? Consider whether it might depend on an
757 artifact of that filter or analysis. Do your conclusions involve temporal or causal
758 dependencies between events in the EEG and events in the world? Make sure that
759 you fully understand how they might be affected by filtering or the TF analysis.

760 **Go with the zeitgeist.** This is in counterpoint to the previous recommendations.
761 One cannot ignore that many studies, past and present, employ filters in ways
762 that we describe as problematic. Those results cannot be discarded, and one may
763 need to use similar methods oneself to allow comparisons, and place new results
764 within the context of prior knowledge. Many researchers and laboratories have
765 well established methodologies that may need to be adhered to for consistency.
766 If such is the case, go for it, but don't forget to fully document, and do call the
767 reader's attention to potential issues.

768 **Conclusion**

769 Filters are ubiquitous in electrophysiology and neuroscience and are an important
770 part of the methodology of any study. Their role is to suppress noise and enhance
771 target activity, but they may have deleterious effects that the investigator should
772 be aware of. When reporting results, it is important to provide enough details so
773 that the reader too can be aware of these potential effects. In some cases there
774 exist alternatives to filtering that are worth considering, in others a filter cannot be

775 avoided. In every case, care must taken to fully understand and report the potential
776 effects of filtering on the patterns reported.

777 **Acknowledgements**

778 The first author was supported by the EU H2020-ICT grant 644732 (COCOHA),
779 and grants ANR-10-LABX-0087 IEC and ANR-10-IDEX-0001-02 PSL*. The
780 second author was supported by grant 390/12 from the Israel Science Foundation
781 (ISF) and by ERC advanced grant GA-340063 (project RATLAND). Giovanni Di
782 Liberto offered useful comments on previous manuscripts.

783 **References**

- 784 Acunzo DJ, MacKenzie G, van Rossum MCW (2012) Systematic biases in early
785 ERP and ERF components as a result of high-pass filtering. *Journal of Neuro-*
786 *science Methods* 209:212–218.
- 787 Bigdely-Shamlo N, Mullen T, Kothe C, Su KM, Robbins KA (2015) The PREP
788 pipeline: standardized preprocessing for large-scale EEG analysis. *Frontiers in*
789 *Neuroinformatics* 9:16.
- 790 Buzsáki G (2006) *Rhythms Of The Brain* Oxford Univ. Press.
- 791 Cohen MX (2017) Rigor and replication in time-frequency analyses of cognitive
792 electrophysiology data. *International Journal of Psychophysiology* 111:80–87.
- 793 Cohen MX (2014) *Analyzing Neural Time Series Data - Theory and practice* MIT
794 Press.
- 795 Cole SR, Voytek B (2017) Brain Oscillations and the Importance of Waveform
796 Shape. *Trends in Cognitive Sciences* 21:137–149.
- 797 Crosse MJ, Di Liberto GM, Bednar A, Lalor EC (2016) The Multivariate Tem-
798 poral Response Function (mTRF) Toolbox: A MATLAB Toolbox for Re-
799 lating Neural Signals to Continuous Stimuli. *Frontiers in Human Neuro-*
800 *science* 10:1–14.

- 801 da Silva FL (2013) EEG and MEG: Relevance to Neuroscience. *Neu-*
802 *ron* 80:1112–1128.
- 803 de Cheveigné A (2016) Sparse Time Artifact Removal. *Journal of Neuroscience*
804 *Methods* 262:14–20.
- 805 de Cheveigné A, Arzounian D (2015) Scanning for oscillations. *Journal of Neural*
806 *Engineering* 12:66020.
- 807 de Cheveigné A, Arzounian D (2018) Robust detrending, rereferencing, outlier
808 detection, and inpainting of multichannel data. *NeuroImage* 172:903–912.
- 809 de Cheveigné A, Parra LC (2014) Joint decorrelation, a versatile tool for multi-
810 channel data analysis. *NeuroImage* 98:487–505.
- 811 de Cheveigné A, Simon JZ (2007) Denoising based on time-shift PCA. *Journal*
812 *of Neuroscience Methods* 165:297–305.
- 813 Delorme A, Palmer J, Onton J, Oostenveld R, Makeig S (2012) Independent EEG
814 sources are dipolar. *PLoS ONE* 7:e30135.
- 815 Ding N, Simon JZ (2013) Adaptive Temporal Encoding Leads to a
816 Background-Insensitive Cortical Representation of Speech. *Journal of Neu-*
817 *roscience* 33:5728–5735.
- 818 Fries P, Scheeringa R, Oostenveld R (2008) Finding Gamma. *Neuron* 58:303–305.
- 819 Gross J, Baillet S, Barnes GR, Henson RN, Hillebrand A, Jensen O, Jerbi K,
820 Litvak V, Maess B, Oostenveld R, Parkkonen L, Taylor JR, van Wassenhove V,
821 Wibral M, Schoffelen JM (2013) Good practice for conducting and reporting
822 MEG research. *NeuroImage* 65:349–363.
- 823 Huigen E, Peper a, Grimbergen Ca (2002) Investigation into the origin of the
824 noise of surface electrodes. *Medical & biological engineering & comput-*
825 *ing* 40:332–338.
- 826 Jones SR (2016) When brain rhythms aren't rhythmic': implication for their
827 mechanisms and meaning. *Current Opinion in Neurobiology* 40:72–80.

- 828 Kappenman ES, Luck SJ (2010) The effects of electrode impedance on
829 data quality and statistical significance in ERP recordings. *Psychophysiol-*
830 *ogy* 47:888–904.
- 831 Keil A, Debener S, Gratton G, Junghöfer M, Kappenman ES, Luck SJ, Luu P,
832 Miller GA, Yee CM (2014) Committee report: Publication guidelines and
833 recommendations for studies using electroencephalography and magnetoen-
834 cephalography. *Psychophysiology* 51:1–21.
- 835 Kıraç L, Vollmar C, Rémi J, Loesch A, Noachtar S (2015) Notch filter arte-
836 fact mimicking high frequency oscillation in epilepsy. *Clinical Neurophysiol-*
837 *ogy* 127.
- 838 Kriegeskorte N, Simmons WK, Bellgowan PSF, Baker CI (2009) Circular anal-
839 ysis in systems neuroscience: the dangers of double dipping. *Nature Neuro-*
840 *science* 12:535–540.
- 841 Lalor EC, Foxe JJ (2010) Neural responses to uninterrupted natural speech can
842 be extracted with precise temporal resolution. *European Journal of Neuro-*
843 *science* 31:189–193.
- 844 Lammertmann C, Lütkenhöner B (2001) Near-DC magnetic fields following
845 a periodic presentation of long-duration tonebursts. *Clinical Neurophysiol-*
846 *ogy* 112:499–513.
- 847 Lopez-Calderon J, Luck SJ (2014) ERPLAB: an open-source toolbox for the
848 analysis of event-related potentials. *Frontiers in Human Neuroscience* 8:1–14.
- 849 Lozano-Soldevilla D (2018) Nonsinusoidal neuronal oscillations: bug or feature?
850 *Journal of Neurophysiology* 119:1595–1598.
- 851 Luck S (2014) *An introduction to the event-related potential technique* MIT press.
- 852 Meyer L (2017) The neural oscillations of speech processing and language com-
853 prehension: State of the art and emerging mechanisms. *European Journal of*
854 *Neuroscience* 48:2609–2621.
- 855 Molden S, Moldestad O, Storm JF (2013) Estimating extracellular spike wave-
856 forms from CA1 pyramidal cells with multichannel electrodes. *PLoS ONE* 8.

- 857 Nikulin VV, Nolte G, Curio G (2011) A novel method for reliable and fast extrac-
858 tion of neuronal EEG/MEG oscillations on the basis of spatio-spectral decom-
859 position. *NeuroImage* 55:1528–1535.
- 860 Nunes P, Srinivasan R (2006) *Electric Fields of the Brain: The Neurophysics of*
861 *EEG* Oxford University Press.
- 862 O’Sullivan JA, Shamma SA, Lalor EC (2015) Evidence for Neural Computations
863 of Temporal Coherence in an Auditory Scene and Their Enhancement during
864 Active Listening. *Journal of Neuroscience* 35:7256–7263.
- 865 Oswal A, Jha A, Neal S, Reid A, Bradbury D, Aston P, Limousin P, Foltynie
866 T, Zrinzo L, Brown P, Litvak V (2016) Analysis of simultaneous MEG and
867 intracranial LFP recordings during Deep Brain Stimulation: A protocol and
868 experimental validation. *Journal of Neuroscience Methods* 261:29–46.
- 869 Parra LC, Spence CD, Gerson AD, Sajda P (2005) Recipes for the linear analysis
870 of EEG. *Neuroimage* 28:326–341.
- 871 Picton TW, Bentin S, Berg P, Donchin E, Hillyard SA, Johnson R, Miller GA, Rit-
872 ter W, Ruchkin DS, Rugg MD, Taylor MJ (2000) Guidelines for using human
873 event-related potentials to study cognition: recording standards and publication
874 criteria. *Psychophysiology* 37:127–152.
- 875 Picton TW, Woods DL, Proulx GB (1978) Human auditory sustained potentials. I.
876 The nature of the response. *Electroencephalography and Clinical Neurophysi-*
877 *ology* 45:186–197.
- 878 Puce A, Hämäläinen MS (2017) A review of issues related to data acquisition and
879 analysis in EEG/MEG studies. *Brain Sciences* 7.
- 880 Quian Quiroga R (2009) What is the real shape of extracellular spikes? *Journal*
881 *of Neuroscience Methods* 177:194–198.
- 882 Rousset GA (2012) Does filtering preclude us from studying ERP time-courses?
883 *Frontiers in Psychology* 3:1–9.
- 884 Saleem AB, Lien AD, Krumin M, Haider B, Rosón MR, Ayaz A, Reinhold K,
885 Busse L, Carandini M, Harris KD, Carandini M (2017) Subcortical Source and

- 886 Modulation of the Narrowband Gamma Oscillation in Mouse Visual Cortex.
887 *Neuron* 93:315–322.
- 888 Schmolesky MT, Wang Y, Hanes DP, Thompson KG, Leutgeb S, Schall JD, Lev-
889 enthal AG (1998) Signal Timing Across the Macaque Visual System. *Journal*
890 *of Neurophysiology* 79:3272–3278.
- 891 Singer W (2018) Neuronal oscillations: Unavoidable and useful? *European*
892 *Journal of Neuroscience* 48:2389–2398.
- 893 Southwell R, Baumann A, Gal C, Barascud N, Friston K, Chait M (2017) Is
894 predictability salient? a study of attentional capture by auditory patterns.
895 *Philosophical Transactions of the Royal Society of London B: Biological Sci-*
896 *ences* 372:doi:10.1098/rstb.2016.0105.
- 897 Tanner D, Morgan-Short K, Luck SJ (2015) How inappropriate high-pass filters
898 can produce artifactual effects and incorrect conclusions in ERP studies of lan-
899 guage and cognition. *Psychophysiology* 52:997–1009.
- 900 Tanner D, Norton JJ, Morgan-Short K, Luck SJ (2016) On high-pass filter artifacts
901 (they’re real) and baseline correction (it’s a good idea) in ERP/ERMF analysis.
902 *Journal of Neuroscience Methods* 266:166–170.
- 903 van Ede F, Quinn AJ, Woolrich MW, Nobre AC (2018) Neural Oscilla-
904 tions: Sustained Rhythms or Transient Burst-Events? *Trends in Neuro-*
905 *sciences* 41:415–417.
- 906 Vanhatalo S, Voipio J, Kaila K (2005) Full-band EEG (FbEEG): An emerging
907 standard in electroencephalography. *Clinical Neurophysiology* 116:1–8.
- 908 VanRullen R (2011) Four common conceptual fallacies in mapping the time
909 course of recognition. *Frontiers in Psychology* 2:1–6.
- 910 Vrba J, Robinson SE (2001) Signal Processing in Magnetoencephalography.
911 *Methods* 25:249–271.
- 912 Widmann A, Schröger E (2012) Filter effects and filter artifacts in the analysis of
913 electrophysiological data. *Frontiers in Psychology* 3:1–5.

- 914 Widmann A, Schröger E, Maess B (2015) Digital filter design for electro-
915 physiological data - a practical approach. *Journal of Neuroscience Meth-*
916 *ods* 250:34–46.
- 917 Yael D, Bar-Gad I (2017) Filter based phase distortions in extracellular spikes.
918 *PLoS ONE* 12:1–13.
- 919 Yeung N, Bogacz R, Holroyd CB, Cohen JD (2004) Detection of synchronized os-
920 cillations in the electroencephalogram: An evaluation of methods. *Psychophys-*
921 *iology* 41:822–832.
- 922 Yuval-Greenberg S, Tomer O, Keren AS, Nelken I, Deouell LY (2008) Transient
923 Induced Gamma-Band Response in EEG as a Manifestation of Miniature Sac-
924 cades. *Neuron* 58:429–441.
- 925 Zoefel B, Heil P (2012) Detection of Near-Threshold Sounds is Independent of
926 EEG Phase in Common Frequency Bands. *Frontiers in Psychology* 4:262.
- 927 Zoefel B, VanRullen R (2017) Oscillatory mechanisms of stimulus processing
928 and selection in the visual and auditory systems: State-of-the-art, speculations
929 and suggestions. *Frontiers in Neuroscience* 11:1–13.